DEVELOPMENTS IN DIGITAL SIGNAL PROCESSING FOR HANDHELD ELECTROACOUSTIC INSTRUMENTATION

P Darlington (1), R G Tyler (2)

(1) University of Salford, Department of Applied Acoustics, Salford M5 4WT
(2) Lucas CEL Instruments Ltd., 35/37 Bury Mead Rd., Hitchin SG5 1RT

1.0 INTRODUCTION

Sound level meters ("Instruments for the measurement of certain frequency and time weighted sound pressure levels", [1]) have evolved from clumsy bulky instruments in the 1930's to the compact handheld devices used today. This process of miniaturization was made possible by developments in solid state electronics which have also enabled a significant increase in performance.

Modern sound level meters are now widely using digital electronics, particularly in the user interface and low bandwidth post-processing functions, which allows a further increase in miniaturisation, reliability and ruggedness of the instrument and ease of manufacture.

Advances in the subject of Digital Signal Processing ("DSP") and, particularly, the introduction of microprocessor chips optimised to perform the kernel functions of DSP, now make an all-digital handheld sound level meter conceivable. In such an instrument, the signal from the microphone is converted to a discrete-time, quantised amplitude representation immediately after the pre-amplifier, with all the weighting and averaging functions implemented by digital signal processing algorithms. This paper defines the computational requirements of such a sound level meter and demonstrates that modern DSP-specific microprocessors and their associated data converters can achieve the required performance at a cost and power consumption within the range possible in a commercially viable handheld sound level meter.

2.0 DIGITAL REPRESENTATION OF THE PRESSURE SIGNAL

In the proposed all-digital sound level meter the output from a conventional instrumentation grade microphone would be converted to a discrete time, quantised amplitude representation, such that the weighting and averaging functions of a sound level meter can be implemented by digital filtering techniques. Such a "sampled" representation of the pressure signal detected by the microphone introduces a bandwidth constraint and an observation noise / amplitude resolution limit. By comparing these two phenomena with the performance required of a modern sound level meter in sections 2.1 and 2.2, it is possible to define a minimum specification for the data converter used to sample the microphone signal, as presented in section 2.3. Note that the performance required of the data converter would also be that required of a "direct-digital-output" microphone, should such a device be developed for instrumentation grade applications.

2.1 Bandwidth Limitations imposed by Discrete-time Operation

If the sampling operation on the microphone signal is performed at a frequency of f_s Hz, then Shannon's theorem states that the data applied to the converter must be lowpass limited to $f_s/2$ to avoid the nonlinear error effects introduced by aliasing. The all-digital sound level meter will have, therefore, a bandwidth of less than 0 -> $f_s/2$ Hz (the actual cut-off frequency of the lowpass anti-aliasing filter will be less than $f_s/2$ as the finite roll-off rate of any practical filter must ensure that the data input at or above $f_s/2$ is attenuated to below the noise floor of the instrument).

The frequency range required of a sound level meter by current British (and related international) Standards [1] is limited at the upper end by the tolerances placed upon the frequency weighting networks to:

type 0 .... $f_{upper} > 20000$ Hz

type 1 .... $f_{upper} > 12500$ Hz

type 2 .... $f_{upper} > 8000$ Hz

type 3 .... $f_{upper} > 8000$ Hz

As any new high performance "all-digital" device would be aimed towards the higher end of the market, where current analog based meters achieve "type 0" bandwidths (even if sold as type 1 instruments) this dictates that for a viable all-digital sound level meter:

$$f_{upper} > 20000 \text{ Hz}$$

which, allowing for the roll-off rate of practical anti-aliasing systems, dictates that:

$$f_s > 50000 \text{ Hz}$$

If the instrument is required to perform 1/nth octave or other narrow-band analysis, it is not unreasonable to expect the meter to be able to record up to the upper end of the 16kHz octave band (approximately 22.6 kHz). This suggests a higher sample frequency of :

$$f_s > 60000 \text{ Hz.}$$

Note that increasing the sample rate not only requires a faster data converter, but also increases the information rate to the digital signal processing chips, increasing the computational load (see section 3).

2.2 Amplitude Quantisation.
The output of an analog to digital converter system is a number stream, representing the input waveform sampled every f_s^{-1} seconds and quantised into one of N quantisation steps. As the converter is built around binary logic, the number of quantising steps generally takes the form:

$$N = 2^n$$

where n is an integer representing the number of binary digits into which the output is coded.

The input analog waveform can assume a range of values and still be quantised to the same quantisation level; the process of amplitude quantisation involves, therefore, the corruption of the data by the addition of a noise component, called quantisation noise. The consequence of the quantisation noise, in the context of the proposed all-digital sound level meter, is to limit the dynamic range of the instrument, as is described below.

Conventional simplified analyses of the influence of converter quantisation noise assume a distribution for the signal and a distribution for the quantising noise and calculate a dynamic range of the data converter defined as a ratio of maximum signal variance to quantisation noise variance. Typical predictions of converter dynamic range [2] are:

$$\text{Dynamic Range} = 6n - 1.24 \ \text{(dB)} \qquad (2)$$

this equation gives rather optimistic predictions.

The dynamic range of a sound level meter can be more usefully described in terms of the ratio of full scale signal to the point at which the instrument fails the 0.4 dB tolerance limit for the maximum error in a change in amplitude of 10 dB [1]. If the quantisation noise has variance σ^2_{noise} and the signal has variance σ^2_{signal}, the meter will fail this tolerance if:

$$10 \log_{10} (\sigma^2_{signal} + \sigma^2_{noise}) - 10 \log_{10} (\sigma^2_{signal}) > 0.4$$

This means that the lowest possible signal, defining the bottom of the dynamic range of the instrument, must have over ten times the variance of the quantising noise :

$$10 \log_{10} (\sigma^2_{min.signal} / \sigma^2_{noise}) > 10.15 \text{ dB} \qquad (3)$$

The dynamic range of the converter system of an all-digital sound level meter to the standards defined in [1] can never exceed the optimistic prediction obtained by combining Equations (2) and (3):

$$\text{Dynamic range} < 6n - 1.24 - 10.15 \text{ dB} \qquad (4)$$

If 60 dB is assumed to be the smallest commercially viable dynamic range available in one range (a wider range of input sound pressures can be transduced using switched gains in the analog pre-amplifier) then the number of converter bits can be deduced from Equation (4) :

$$\begin{aligned} n &= (60 + 1.24 + 10.15) / 6 \qquad (5) \\ &= 12 \ \text{(n integer)} \end{aligned}$$

The authors have conducted extensive tests on a contemporary "industry standard" 16 bit (n=16) sigma-delta converter and have achieved a dynamic range of between 74 and 78 dB, significantly lower than the 84.6 dB range predicted by Equation (4), as a result of the underestimate of the quantisation noise embodied in Equation (2).

2.3 Minimum Specification of the Data Converter in a Digital Sound Level Meter.
Given the constraints introduced by time domain sampling and amplitude domain quantisation, discussed in the sections above, it is possible to define minimum

requirements which the analog to digital converter system at the front end of an all-digital sound level meter must achieve (rather than being marketed as *capable* of achieving) as :

sample frequency	(f_s)	> 60 kHz
resolution	(n)	> 12 bits

The signal processing electronics after the data converter may have finite amplitude resolution and dynamic range, in which case the performance of the entire instrument may be worse than that of the front end data converter alone - it is therefore emphasized that the specification above is for an absolute minimum standard which should be exceeded if at all possible.

3.0 SIGNAL PROCESSING REQUIREMENTS FOR DIGITAL WEIGHTING AND AVERAGING FILTERS

Having demonstrated in the previous section that the resolution and bandwidth of contemporary analog to digital converters meets the requirements for the data converter in a sound level meter, it is appropriate to consider the computational load required to perform the weighting and averaging functions required of a sound level meter to current standards. The implementation of integrating sound level meter functions will not be considered as the computations involved are performed at low bandwidth and are often implemented in software in current meters.

In order to establish the computational power required to implement the sound level meters functions, the algorithms used for frequency weighting and averaging will be examined in the following sections.

3.1 Digital Implementation of Frequency Weighting Networks.

The A and C frequency weighting networks, specified in [1], will be required in any commercially viable sound level meter and can be implemented using digital filtering algorithms running on digital signal processing specific microprocessors within the instrument. These processors are generally equipped with an instruction set optimised to perform the math-intensive operations involved in digital filtering, often by the on-chip provision of a sophisticated arithmetic logic unit, usually featuring a hardware multiplier. The speed and programming flexibility of the device is usually further enhanced with complex bus structures, supporting advanced parallel addressing modes and some degree of instruction pipelining. The weighting networks (and a high pass filter defining the bottom of the bandwidth of the instrument) can be implemented on the devices using standard infinite impulse response digital filter structures.

The frequency weighting characteristics can be realised as digital filters by taking the s-plane pole/zero locations specified in [1] and transforming these to z-plane digital filter coefficients using the "impulse invariant" technique [2]. This relates a pole location in the Laplace plane, for example the poles in the C weighting associated with the roll-off at 20.6 Hz (s = -129.4 rad/sec), to the equivalent z domain pole at :

$$z_p = e^{-(129.4/fs)}$$

which for a sample rate of 60 kHz gives a pole at:

$$z_p = 0.997845....... + j0$$

The proximity of the low frequency poles to the unit circle and the precision to which the pole must be specified can cause problems in the context of the finite precision number representation schemes used in DSP chips - this is discussed below.

Using the impulse invariant design technique it has been found possible to implement the digital weighting filters using the same number of poles and zeros as specified for the analog versions, plus an additional zero for high frequency trimming. The number of singularities required to implement the weighting filters and a high pass d.c. reject filter, for the entire digital sound level meter is :

Filter	No. of Poles	No. of (non-zero) Zeros
High Pass	1	1
C Weight	4	2
A weight	2	2
HF Trimming	0	1
TOTAL	7	6

Since both C and A weightings are built from real pole/zero locations, the digital weighting filters can be realised as a cascade of seven bilinear sections, each of form:

$$H(z) = S (z - b_1)(z - a_1)^{-1} \qquad (6)$$

where z is the unit advance operator, H is the response of the section, S is a gain term for the section and a_1 and b_1 are the pole and zero locations respectively. There are two very considerable advantages in using cascaded bilinear sections as opposed to other more compact higher order forms ;

a) errors in coefficient locations due to the finite precision of a fixed point number representation system (or a fixed length mantissa in a floating point context) do not compound under multiplication in the calculation of z-domain polynomials,

and

b) the dynamic range of the data variables is maximised (the dynamic range of the system can be eroded by the requirement for large headroom for intermediate results, particularly in the context of compacted "lattice" filter structures).

As the difference equation equivalent of Equation (6):

$$y_k = S(x_k - b_1x_{k-1}) + a_1y_{k-1}$$

(where y_k and x_k are the section input and output at time k respectively) requires 3 multiplications and two additions, the total arithmetic load on the processor implementing the filtering functions in a digital sound level meter is:

$21 \times f_s$ multiplications per second

$14 \times f_s$ additions per second

which assuming floating point operation (see section 3.2) and a sample rate of 60 kHz, is a computational load of 2.1×10^6 floating point operations per second (2.1 MFLOPS).

3.2 Averaging Filters

Having computed the weighted sound pressures, it is necessary to square and average the instantaneous signals. The squaring (which must be applied three times per period to calculate instantaneous values of the squared linear, C weighted and A weighted pressures) amounts to a multiplication operation and so adds a further 3 x f_s arithmetic operations per second to the computational load.

In performing the squaring operation, the required dynamic range of the processor must be double that of the converter if the dynamic range of the quantised signal is not to be degraded. This is not usually a problem as the processor would typically have a data bus width equal to or greater than n (the resolution of the converter) such that the output register of the hardware multiplier would have greater than 2n resolution. If, however, the instantaneous squared pressure signals are to be transmitted, stored or averaged, problems may occur if the processor cannot handle a dynamic range *exceeding* twice the dynamic range of the converter system. *This will practically mean that the processing devices in all-digital sound level meters should support some kind of floating point number representation scheme.*

The averaging of the instantaneous squared (weighted) pressure signals can be performed in two ways, both of which are detailed below.

3.2.1 Energy Accumulation. If the squared (weighted) pressure signals are simply summed over some defined period, the accumulated sum will define a measure of energy which is directly useful in calculating equivalent sound pressure levels, or doses in a "dosemeter" application.

If the summation period chosen extends over L samples, then the number representation scheme in the processor must exceed the dynamic range of the input data by a factor of 2 x L if the meter is to avoid saturation. At the end of every L samples, the dose can be output at the reduced rate of f_s/L to an algorithm which either computes equivalent energy levels (L_{eq}'s) or which *estimates* the sound pressure level. Note that the calculation of sound pressures from accumulated energies will give a different result than that obtained by averaging the instantaneous signals in a first order lowpass filter having time constant appropriate to one of the standard averaging modes in [1].

The process of accumulation of energies introduces an extra full bandwidth arithmetic load on the digital processing device of at least 3 x f_s additions (some multiplication may be required for scaling in a limited resolution environment, although this may be achieved by logical shifts rather than multiplications).

3.2.2 Averaging Instantaneous Pressure Signals by Filtering. A more conventional approach to the calculation of averaged values of the raw squared pressure signals is to pass the data through low pass filter sections. The time constants of these sections could be chosen to correspond to the standard averaging responses [1], although the pole locations involved would be very near the unit circle (and cannot, for example, be encoded in a 16 bit binary number to sufficient accuracy). For this reason, it would be more appropriate to apply averaging filters having a higher frequency cut-off frequency, to downsample to a rate of f_s/L and to subsequently apply the "Impulse, Fast or Slow" characteristic at the lower sample rate. Experiments have demonstrated that this "two-stage" averaging can yield results which do not exceed the tolerancing limits required of the instrument *provided that the full bandwidth averaging filter time constants and the downsample factor L are chosen appropriately.*

Calculating the data sets typical of integrating sound level meters can be achieved from the filtered signal described above using the same techniques as employed in current meters with analog front ends.

The computational cost of averaging the raw squared pressure data by filtering is that associated with a bilinear filter section (for each of Lin. A and C pressure signals):

$3 \times 3 \times f_s$ multiplies per second
$3 \times 2 \times f_s$ additions per second.

4.0 SIGNAL PROCESSING LOAD IN A DIGITAL SOUND LEVEL METER

The total computational load involved in performing the basic full bandwidth "front-end" functions of a sound level meter using digital signal processing techniques can now be calculated, assuming a 60 kHz sample rate and floating point operation:

Averaging Method	Computational Load
Energy Accumulation	2.46 MFLOPS
Filtering	3.18 MFLOPS

Although this computational load is well within the capability of several DSP specific microprocessors, (for example 25 MFLOPS max. for the AT&T DSP32C) the performance must be achieved in the context of a handheld instrument, which must compete with the high performance analog designs currently in the marketplace.

Both power consumption and the prime cost of the components strongly favour the analog solution, as the high performance DSP processors are both power hungry and expensive. Current LUCAS CEL analog sound level meters can achieve the frequency weighting and time averaging functions using components which consume only 8% of the total power used by the instrument and account for only 6% of the prime component cost. Feasibility studies have identified a DSP chip which could perform the weighting and averaging functions, using the algorithms described in the previous sections, but which would consume 20% of the instrument's power and account for 15% of the total cost. These costs represent a disadvantage in approaching the front end processing requirements of a sound level meter using digital filtering techniques, but the advantages of repeatability of manufacture, standardization of performance and scope for future product enhancements (addition of narrowband analysis capability, etc.) could more than compensate for these drawbacks.

5.0 CONCLUDING REMARKS

This paper has demonstrated that the front end full bandwidth signal processing functions required in a sound level meter could be practically implemented using digital signal processing techniques, running on currently available electronics. The resolution and bandwidth with which the input data must be acquired has been demonstrated to be within the capacity of modern data converter systems. The mathematically intensive operations involved in forming the required time and frequency weightings have been used to calculate the computational load in the DSP processor.

Although an all digital sound level meter has some drawbacks (dynamic range, power consumption and price) the authors believe that the considerable advantages which the digital technology offers to both manufacturers and users in reliability, enhanced performance and accuracy will ensure an increasing use of Digital Signal Processing in Handheld Electroacoustic Instrumentation in coming years.

REFERENCES

[1] British Standards Institution, "*Specification for Sound Level Meters*" BS 5969:1981, IEC 651:1979.

[2] A V Oppenheim & R W Schafer, "*Digital Signal Processing*", Prentice Hall 1975